arXiv:1402.5715v2 [stat.ML] 1 Mar 2014

Variational Particle Approximations

Tejas D. Kulkarni*[1,2] **Ardavan Saeedi***[2] **Samuel Gershman**[1]
[1] Department of Brain & Cognitive Sciences, MIT
[2] Computer Science & Artificial Intelligence Laboratory, MIT
* These authors contributed equally to this work and listed alphabetically.

Abstract

Monte Carlo methods provide a powerful framework for approximating probability distributions with a set of stochastically sampled particles. In this paper, we rethink particle approximations from the perspective of variational inference, where the particles play the role of variational parameters. This leads to a deterministic version of Monte Carlo in which the particles are selected to optimize the Kullback-Leibler divergence between the approximation and the target distribution. Variational particle approximations overcome some of the weaknesses of Monte Carlo methods like particle filtering, leading to substantially improved performance on several synthetic and real-world datasets.

1 Introduction

Monte Carlo methods are based on the idea that one can approximate a complex distribution with a set of stochastically sampled particles. The flexibility and variety of Monte Carlo methods have made them the workhorse of statistical computation. However, their success relies critically on having available a good sampler, and designing such a sampler is often challenging.

In this paper, we rethink particle approximations from a different perspective. Suppose we got to pick where to place the particles; where would we put them? Intuitively, we would want to distribute them in such a way that they cover high probability regions of the target distribution, but without the particles all devolving onto the mode of the distribution. This problem can be formulated precisely within the framework of variational inference (Wainwright and Jordan, 2008), which treats probabilistic inference as an optimization problem over a set of distributions. We derive a coordinate ascent update for particle approximations that iteratively minimizes the Kullback-Leibler (KL) divergence between the particle approximation and the target distribution.

After introducing our general framework, we describe how it can be applied to filtering and smoothing problems. We then show experimentally that variational particle approximations can overcome a number of problems that vex conventional Monte Carlo methods. In particular, our approach is able to produce a diverse, high probability set of particles in situations where Monte Carlo methods sometimes degenerate.

2 Background

Consider the problem of approximating a probability distribution $P(x)$ over discrete latent variables $x = \{x_1, \ldots, x_N\}, x_n \in \{1, \ldots, M_n\}$, where the target distribution is known only up to a normalizing constant Z: $P(x) = f(x)/Z$. We further assume that $P(x)$ is a Markov network defined on a graph G, so that $f(x)$ factorizes according to:

$$f(x) = \prod_c f_c(x_c), \tag{1}$$

where $c \subseteq \{1, \ldots, N\}$ indexes the maximal cliques of G.

2.1 Importance sampling and sequential Monte Carlo

A general way to approximate $P(x)$ is with a weighted collection of K particles, $\{x^1, \ldots, x^K\}$:

$$P(x) \approx Q(x) = \sum_{k=1}^{K} w^k \delta[x, x^k], \tag{2}$$

where $\delta[\cdot,\cdot] = 1$ if its arguments are equal and 0 otherwise. Importance sampling is a Monte Carlo method that stochastically generates particles from a proposal distribution, $x^k \sim \phi(\cdot)$, and computes the weight according to $w^k \propto f(x^k)/\phi(x^k)$. Importance sampling has the important property that the particle approximation converges to the target distribution as $K \to \infty$ (Robert and Casella, 2004).

Sequential Monte Carlo methods such as particle filtering (Doucet et al., 2001) apply importance sampling to stochastic dynamical systems (where n indexes time) by sequentially sampling the latent variables at each time point using a proposal distribution $\phi(x_n|x_{n-1})$. This procedure can produce conditionally low probability particles; therefore, most algorithms include a resampling step which replicates high probability particles and kills off low probability particles. The downside of resampling is that it can produce degeneracy: the particles become concentrated on a small number of hypotheses, and consequently the effective number of particles is low.

2.2 Variational inference

Variational methods (Wainwright and Jordan, 2008) define a parametrized family of probability distributions $\mathcal{Q}$ and then choose $Q \in \mathcal{Q}$ that maximizes the *variational free energy*:

$$\mathcal{L}[Q] = \sum_x Q(x) \log \frac{f(x)}{Q(x)}. \tag{3}$$

The variational free energy is related to the normalizing constant (partition function) Z and the KL divergence through the following identity:

$$\log Z = \mathrm{KL}[Q||P] + \mathcal{L}[Q], \tag{4}$$

where

$$\mathrm{KL}[Q||P] = \sum_x Q(x) \log \frac{Q(x)}{P(x)}. \tag{5}$$

Since $\mathrm{KL}[Q||P] \geq 0$, the variational free energy is a lower bound on the log partition function, achieving equality when the KL divergence is minimized to 0. Maximizing $\mathcal{L}[Q]$ with respect to Q is thus equivalent to minimizing the KL divergence between Q and P.

Unlike the Monte Carlo methods described in the previous section, variational methods do not in general converge to the target distribution, since typically P is not in $\mathcal{Q}$. The advantage of variational methods is that they guarantee an improved bound after each iteration, and convergence is easy to monitor (unlike most Monte Carlo methods). In practice, variational methods are also often more computationally efficient.

We next consider particle approximations from the perspective of variational inference. We then turn to the application of particle approximations to inference in stochastic dynamical systems.

3 Variational particle approximations

Variational inference can be connected to Monte Carlo methods by viewing the particles as a set of variational parameters parameterizing Q. For the particle approximation defined in Eq. 2, the variational free energy takes the following form:

$$\mathcal{L}[Q] = \sum_{k=1}^{K} w^k \log \frac{f(x^k)}{w^k}. \tag{6}$$

Using Lagrange multipliers, it can be shown that maximizing $\mathcal{L}[Q]$ by coordinate ascent can be performed by the following updates for each variable n:

1. For $m = 1, \ldots, M_n$ and $k = 1, \ldots, K$ compute new unnormalized weights:

$$W_n(k, m) = w^k \frac{\mathcal{F}_n(\tilde{x}^k)}{\mathcal{F}_n(x^k)} \tag{7}$$

where $\mathcal{F}_n(x) = \prod_{c:n\in c} f_c(x_c)$ and

$$\tilde{x}_j^k = \begin{cases} x_j^k & \text{if } j \neq n \\ m & \text{if } j = n \end{cases} \tag{8}$$

2. Assemble the K largest elements of W_n and their associated particles into a new particle approximation $\{w^k, x^k\}$, normalizing the weights so that they sum to 1.

We assume that the particles are initialized randomly, and the initial weights are set to $w^k \propto$

$\mathcal{F}_n(\tilde{x}^k)$. We refer to this algorithm as *particle variational inference* (PVI). When repeatedly iterated, PVI will converge to a local maximum of the variational free energy. Note that in principle more sophisticated methods can be used to find the top K modes (e.g., Flerova et al., 2012; Yanover and Weiss, 2003); however, we have found that this coordinate ascent algorithm is fast, easy to implement, and very effective in practice (as our experiments below demonstrate).

An important aspect of this framework is that it maintains the same asymptotic guarantee as importance sampling: Q converges to P as $K \to \infty$. Thus, PVI combines advantages of variational methods (monotonically decreasing KL divergence between Q and P) with the asymptotic correctness of Monte Carlo methods.

4 Filtering and smoothing in hidden Markov models

We now describe how variational particle approximations can be applied to filtering and smoothing in hidden Markov models (HMMs). Consider a hidden Markov model with observations $y = \{y_1, \ldots, y_N\}$ generated by the following stochastic process:

$$P(y, x, \theta) = P(\theta) \prod_n P(y_n|x_n, \theta) P(x_n|x_{n-1}, \theta), \quad (9)$$

where θ is a set of transition and emission parameters. We are particularly interested in *marginalized* HMMs where the parameters are integrated out: $P(y, x) = \int_\theta P(y, x, \theta) d\theta$. This induces dependencies between observation n and all previous observations, making inference challenging.

Filtering is the problem of computing the posterior over the latent variables at time n given the history $y_{1:n}$. To construct the variational particle approximation of the filtering distribution, we need to compute the product of potentials for variable n:

$$\mathcal{F}_n(x) = P(y_n|x_{1:n}, y_{1:n-1}) P(x_n|x_{1:n-1}). \quad (10)$$

We can then apply the coordinate ascent update described in the previous section. This update is simplified in the filtering context due to the underlying Markov structure:

$$W_n(k, m) = w^k P(y_n|x_n^k = m, x_{1:n-1}, y_{1:n-1}) \times P(x_n^k = m|x_{1:n-1}). \quad (11)$$

At each time step, the algorithm selects the K continuations (new variable assignments of the current particle set) that maximize the variational free energy.

Smoothing is the problem of computing the posterior over the latent variables at time n given data from both the past and the future, $y_{1:N}$. The product of potentials is given by:

$$\mathcal{F}_n(x) = P(y_n|x_{1:n}, y_{-n}) P(x_n|x_{-n}), \quad (12)$$

where x_{-n} refers to all the latent variables except x_n (and likewise for y_{-n}). This potential can be plugged into the updates described in the previous section.

To understand PVI applied to filtering problems, it is helpful to contemplate three possible fates for a particle at time n (illustrated in Figure 1):

- **Selection**: A single continuation of particle k has non-zero weight. This can be seen as a deterministic version of particle filtering, where the sampling operation is replaced with a max operation.
- **Splitting**: Multiple continuations of particle k have non-zero weight. In this case, the particle is split into multiple particles at the next iteration.
- **Deletion**: No continuations of particle k have non-zero weight. In this case, the particle is deleted from the particle set.

Similar to particle filtering with resampling, PVI deletes and propagates particles based on their probability. However, as we show later, PVI is able to escape some of the problems associated with resampling.

5 Related work

PVI is a special case of a *mixture mean-field variational approximation* (Jaakkola and Jordan, 1998; Lawrence, 2000):

$$Q(x) = \sum_{k=1}^{K} Q(k) \prod_{n=1}^{N} Q(x_n|k). \quad (13)$$

In PVI, $Q(k) = w^k$ and $Q(x_n|k) = \delta[x_n, x_n^k]$.

When $K = 1$, PVI is equivalent to *iterated conditional modes* (ICM; Besag, 1986), which iteratively maximizes each latent variable conditional on the rest of the variables.

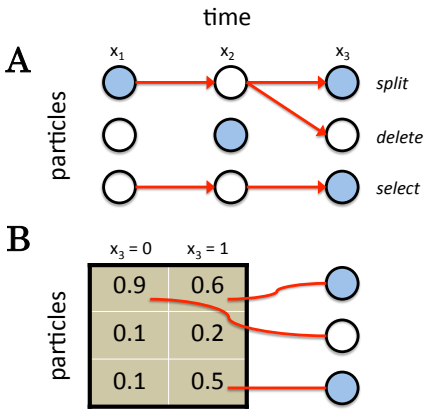

Figure 1: **Schematic of PVI for filtering problems**. (A) Illustration of different filtering scenarios over 3 time steps in a binary state space with $K = 3$ particles. The shading of each circle indicates the binary value of the corresponding variable. Arrows indicate the evolution of the particles. (B). Each cell in the table on the left represents the probability of a particular continuation. The K continuations with highest probability are selected for propagation to the next time step (shown on the right).

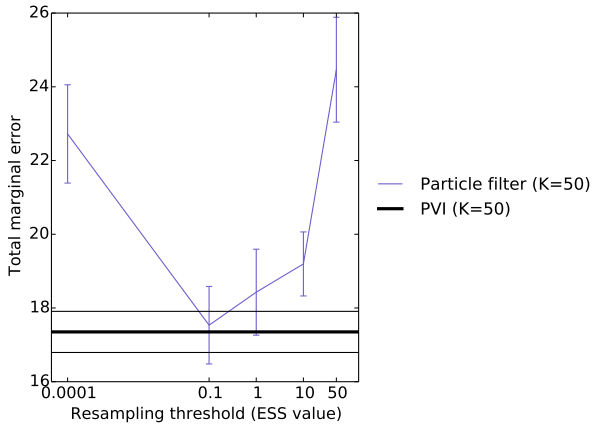

Figure 2: **HMM with binary hidden states and observations.** Total marginal error computed for a sequence of length 200. For particle filtering the total error for every ESS value is averaged over 5 sequences generated from the HMM; in addition, for each sequence we reran the particle filter 5 times (thus 25 runs total). Note the logarithmic scale of the x-axis. Error bars and the thin black lines correspond to standard error of the mean.

PVI is conceptually similar to nonparametric variational inference (Gershman et al., 2012), which approximates the posterior over a continuous state space using a set of particles convolved with a Gaussian kernel.

Frank et al. (2009) used particle approximations within a variational message passing algorithm. The resulting approximation is "local" in the sense that the particles are used to approximate messages passed between nodes in a factor graph, in contrast to the "global" approximation produced by PVI, which attempts to capture the distribution over the entire set of variables.

Ionides (2008) described a truncated version of importance sampling in which weights falling below some threshold are set to 0. When latent variables are sampled exhaustively and without replacement, truncated importance sampling is equivalent to PVI, where the weight threshold is induced by the choice of K.

Finally, our approach is closely related to the problem of finding the K most probable latent variable assignments (Flerova et al., 2012; Yanover and Weiss, 2003). We view this problem through the lens of particle approximations, connecting it to both Monte Carlo and variational methods.

6 Experiments

In this section, we compare the performance of PVI to several widely used approximate inference algorithms, including particle filtering and variational methods. We first present a didactic example to illustrate how PVI can sometimes succeed where particle filtering fails. We then apply PVI to three probabilistic models: the Dirichlet process mixture model (DPMM; Antoniak, 1974; Escobar and West, 1995; Fearnhead, 2004), the infinite HMM (iHMM; Beal et al., 2002; Teh et al., 2006), and the Ising model.

6.1 Didactic example: binary HMM

As a didactic example, we use a simple HMM with binary hidden states (x) and observations (y):

$$\begin{aligned} P(x_{n+1} = 0|x_n = 0) &= \alpha_0 \\ P(x_{n+1} = 1|x_n = 1) &= \alpha_1 \\ P(y_n = 0|x_n = 0) &= \beta_0 \\ P(y_n = 1|x_n = 1) &= \beta_1, \end{aligned} \tag{14}$$

with α_0, α_1, β_0, and β_1 all less than 0.5. We will use this model to illustrate how PVI differs from particle filtering.

For illustration, we use the following parameters:

$\alpha_0 = 0.2$, $\alpha_1 = 0.1$, $\beta_0 = 0.3$, and $\beta_1 = 0.2$. Suppose you observe a sequence generated from this model. For a sufficiently long sequence, a particle filter with resampling will eventually delete all conditionally unlikely particles, and thus suffer from degeneracy. On the other hand, without resampling the approximation will degrade over time because conditionally unlikely particles are never replaced by better particles. For this reason, it is sometimes suggested that resampling only be performed when the effective sample size (ESS) falls below some threshold. The ESS is calculated as follows:

$$\text{ESS} = \frac{1}{\sum_{k=1}^{K} (w^k)^2}. \tag{15}$$

A low ESS means that most of the weight is being placed on a small number of particles, and hence the approximation may be degenerate (although in some cases this may mean that the target distribution is peaky). We evaluated particle filtering with multinomial resampling on synthetic data generated from the HMM described above. Approximation accuracy was measured by using the forward-backward algorithm to compute the hidden state posterior marginals exactly and then comparing these marginals to the particle approximation. Figure 2 shows performance as a function of ESS threshold, demonstrating that there is a fairly narrow range of thresholds for which performance is good. Thus in practice, successful applications of particle filtering may require computationally expensive tuning of this threshold.

In contrast, PVI achieves performance comparable to the optimal particle filter, but without a tunable threshold. This occurs because PVI uses an implicit threshold that is automatically tuned to the problem. Instead of resampling particles, PVI deletes or propagates particles deterministically based on their relative contribution to the variational bound.

6.2 Dirichlet process mixture model

A DPMM generates data from the following process:

$$G \sim \text{DP}(\alpha, G_0), \qquad \theta_n | G \sim G, \qquad y_n | \theta_n \sim F(\theta_n),$$

where $\alpha \geq 0$ is a concentration parameter and G_0 is a base distribution over the parameter θ_n of the observation distribution $F(y_n|\theta_n)$. Since the Dirichlet process induces clustering of the parameters θ into K distinct values, we can equivalently express this model in terms of a distribution over cluster assignments, $x_n \in \{1, \ldots, C\}$. The distribution over x is given by the Chinese restaurant process (Aldous, 1985):

$$P(x_n = c | x_{1:n-1}) \propto \begin{cases} t_c & \text{if } k \leq C_+ \\ \alpha & \text{if } c = C_+ + 1, \end{cases} \tag{16}$$

where t_c is the number of data points prior to n assigned to cluster c and C_+ is the number of clusters for which $t_c > 0$.

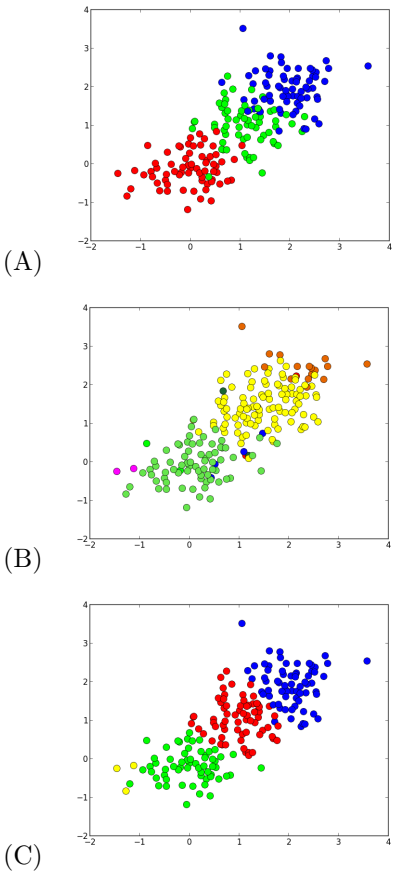

Figure 3: **Synthetic data drawn from a mixture of Gaussians**. (A) Original data, with clusters indicated by colors. We treat this as a filtering problem, analyzing one randomly chosen data point at a time. (B) A representative run of particle filtering. (C) PVI is able to find all three modes and is better able to infer the underlying cluster structure than particle filtering.

6.2.1 Synthetic data

We first demonstrate our approach on synthetic datasets drawn from various mixtures of Gaussians (see Table 1). The model parameters for each of the simulated dataset were set in order to create a spectrum of increasingly overlapping clusters. In particular, we constructed models out of the following

Dataset	Particle filtering ($K = 20$)	PVI ($K = 1$)	PVI ($K = 20$)
D1: $[\mu_1, 4\mu_2, 8\mu_2], \Sigma_1$	0.97±0.03	0.93±0.05	0.99±0.02
D2: $[\mu_1, 4\mu_2, 8\mu_2], \Sigma_2$	0.89±0.05	0.86±0.07	0.90±0.03
D3: $[\mu_1, 2\mu_2, 4\mu_2], \Sigma_1$	0.58±0.12	0.51±0.03	0.74±0.16
D4: $[\mu_1, 2\mu_2, 4\mu_2], \Sigma_2$	0.50±0.06	0.46±0.05	0.55±0.07
D5: $[\mu_1, \mu_2, 2\mu_2], \Sigma_1$	0.05±0.05	0.014±0.02	0.14±0.10
D6: $[\mu_1, \mu_2, 2\mu_2], \Sigma_2$	0.15±0.08	0.11±0.06	0.19±0.07

Table 1: Clustering accuracy (V-Measure) for DPMM. Each dataset consisted of 200 points drawn from a mixture of 3 Gaussians. For each dataset, we repeated the experiment 150 times by iterating through random seeds. The left column shows the ground truth mean for each cluster and the covariance matrix (shared across clusters).

building blocks:

$$\mu_1 = \begin{pmatrix} 0.0, & 0.0 \end{pmatrix}, \qquad \mu_2 = \begin{pmatrix} 0.5, & 0.5 \end{pmatrix}$$
$$\Sigma_1 = \begin{pmatrix} 0.25, & 0.0 \\ 0.0, & 0.25 \end{pmatrix}, \qquad \Sigma_2 = \begin{pmatrix} 0.5, & 0.0 \\ 0.0, & 0.5 \end{pmatrix}.$$

For the DPMM, we used a Normal-Gamma prior with a diagonal covariance matrix. Some example clusterings obtained by different algorithms are shown in Figure 3 for dataset D4.

Clustering accuracy was measured quantitatively using V-measure (Rosenberg and Hirschberg, 2007). As shown in Table 1, we observe only marginal improvements when the means are farthest from each other and variances are small, as these parameters leads to well-separated clusters in the training set. However, the relative accuracy of PVI increases considerably when the clusters are overlapping, either due to the fact that the means are close to each other or the variances are high.

An interesting special case is when $K = 1$. In this case, PVI is equivalent to the greedy algorithm proposed by Daume (2007) and later extended by Wang and Dunson (2011). In fact, this algorithm was independently proposed in cognitive psychology by Anderson (1991). As shown in Table 1, PVI with 20 particles outperforms the greedy algorithm, as well as particle filtering with 20 particles.

6.2.2 Spike sorting

Spike sorting is an important problem in experimental neuroscience settings where researchers collect large amounts of electrophysiological data from multi-channel tetrodes. The goal is to extract from noisy spike recordings attributes such as the number of neurons, and cluster spikes belonging to the same neuron. This problem naturally motivates the use of DPMM, since the number of neurons recorded by a single tetrode is unknown. Previously, Wood and Black (2008) applied the DPMM to spike sorting using particle filtering and Gibbs sampling. Here we show that PVI can outperform particle filtering, achieving high accuracy even with a small number of particles.

We used data collected from a multiunit recording from a human epileptic patient (Quiroga et al., 2004). The raw spike recordings were preprocessed following the procedure proposed by Quiroga et al. (2004), though we note that our inference algorithm is agnostic to the choice of preprocessing. The original data consist of an input vector with 10 dimensions and 9196 data points. We place a Normal-Inverse Wishart prior on the cluster parameters (see Supplementary materials for more details).

We compared our algorithm to the current best particle filtering baseline, which uses stratified resampling (Wood and Black, 2008; Fearnhead, 2004). The same model parameters were used for all comparisons. Qualitative results, shown in Figure 4, demonstrate that PVI is better able to separate the spike waveforms into distinct clusters, despite running PVI with 10 particles and particle filtering with 100 particles. We also provide quantitative results by calculating the held-out log-likelihood on an independent test set of spike waveforms. The quantitative results (summarized in Table 2) demonstrate that even with only 10 particles PVI can outperform particle filtering with 1000 particles.

6.3 Infinite HMM

An iHMM generates data from the following process:

$$G_0 \sim \mathrm{DP}(\gamma, H), \qquad G_k | G_0 \sim \mathrm{DP}(\alpha, G_0),$$
$$x_n | x_{n-1} \sim G_{x_{n-1}}, \qquad \theta_k \sim H, \qquad y_n | x_n \sim F(\theta_{x_n}).$$

Like the DPMM, the iHMM induces a sequence of cluster assignments. The distribution over cluster assignments is given by the Chinese restaurant franchise (Teh et al., 2006). Letting t_{jc} denote the number of times cluster j transitioned to cluster c, x_n is

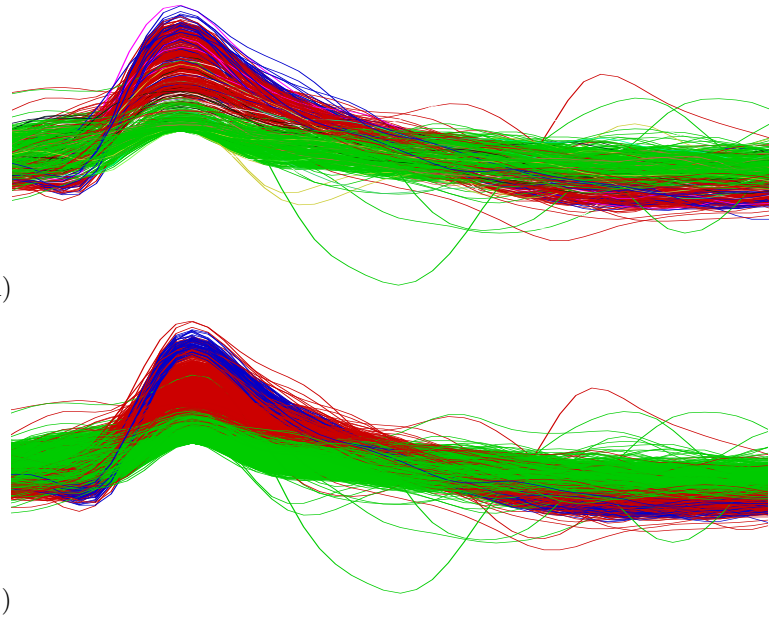

Figure 4: **Spike Sorting using the DPMM**. Each line is an individual spike waveform, colored according to the inferred cluster. (A) Result using particle filtering with 100 particles and stratified resampling as reported in Wood and Black (2008). (B) Result using PVI. The same model parameters were used for both particle filtering and PVI.

assigned to cluster c with probability proportional to $t_{x_{n-1}c}$, or to a cluster never visited from x_{n-1} ($t_{x_{n-1}c} = 0$) with probability proportional to α. If an unvisited cluster is selected, x_n is assigned to cluster c with probability proportional to $\sum_j t_{jc}$, or to a new cluster (i.e., one never visited from any state, $\sum_j t_{jc} = 0$) with probability proportional to γ.

6.3.1 Synthetic data

We generated 50 sequences with length 500 from 50 different HMMs, each with 10 hidden and 5 observed states. For the rows of the transition and initial probability matrices of the HMMs we used a symmetric Dirichlet prior with concentration parameter 0.1; for the emission probability matrix, we used a symmetric Dirichlet prior with concentration parameter 10.

Figure 5A illustrates the performance of PVI and particle filtering (with multinomial and stratified resampling) for varying numbers of particles (1, 10 and 100). Performance error was quantified by computing the Hamming distance between the true hidden sequence and the sampled sequence.[1] The error bars show two standard deviations. The results show that PVI outperforms particle filtering in all three cases.

6.3.2 Text analysis

We next analyzed a real-world dataset, text taken from the beginning of "Alice in Wonderland". We compared PVI to the beam sampler (Van Gael et al., 2008), a combination of dynamic programming and slice sampling, which was previously applied to this dataset. Following Van Gael et al. (2008), the training set consisted of a sequence of 1000 characters from the first chapter of the book. We then calculated the predictive log-likelihood of the subsequent 4000 characters. When computing the predictive log-likelihood, we truncated the number of the states with the maximum number of hidden states sampled in the training phase. For both PVI and particle filtering we fixed the hyperparameters α and γ to 1.

Figure 5B shows the predictive log-likelihood for varying numbers of particles. Even with a small number of particles, PVI can outperform both particle filtering and the beam sampler.

6.4 Ising model

So far, we have been studying inference in directed graphical models, but PVI can also be applied to undirected graphical models. We illus-

[1] The Munkres algorithm was used to maximize the overlap between the two sequences.

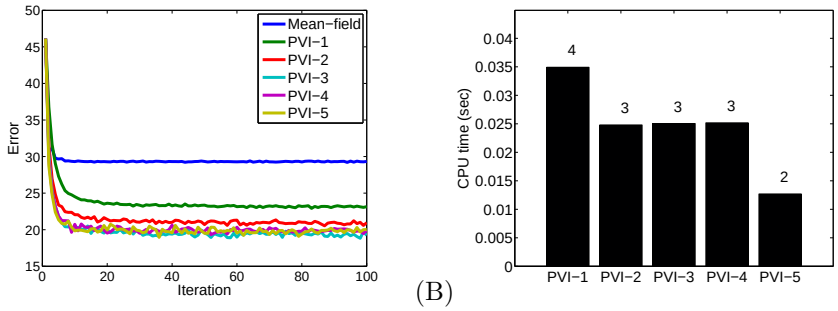

(A) (B)

Figure 6: **Ising model results**. (A) Total variation distance between the inferred and true marginals as a function of iteration. (B). Computation time required to outperform mean-field. The number of iterations is shown above each bar.

Method	**Held-out log-likelihood**
PVI ($K = 10$)	-3.2474×10^5 ($\hat{C} = 3$)
PVI ($K = 100$)	-1.3888×10^5 ($\hat{C} = 3$)
SMC (Stratified) ($K = 10$)	-1.4771±0.21 × 10^6 ($\hat{C} = 37$)
SMC (Stratified) ($K = 100$)	-5.6757±1.14 × 10^5 ($\hat{C} = 13$)
SMC (Stratified) ($K = 1000$)	-3.2965×10^5 ($\hat{C} = 5$)

Table 2: Spike sorting held-out log-likelihood scores for 200 test points. The best performance is achieved by PVI with 100 particles. Shown in parentheses is the *maximum a posteriori* number of clusters, $\hat{C}$.

(a)

Number of particles	**PVI**	**SMC**
$K = 10$	15.20s	14.71s
$K = 50$	153.75s	184.17s
$K = 100$	567.84s	699.43s

(b)

Number of particles	**PVI**	**SMC**
$K = 10$	36.20s	124s
$K = 50$	144.6s	334.2s
$K = 100$	313.8s	454.2s

Table 3: Run time comparison for DPMM. (a) Results using synthetic DPMM dataset from Table 1 and (b) highlights results obtained by using the spike sorting dataset. In both cases, the run time of PVI is slightly better than SMC.

trate this using the Ising model for binary vectors $x \in \{-1, +1\}^N$:

$$f(x) = \frac{1}{2} x W x^\top + \theta x^\top, \qquad (17)$$

where $W \in \mathbb{R}^{N \times N}$ and $\theta \in \mathbb{R}^N$ are fixed parameters. In particular, we study a square-lattice spin-glass, where $W_{ij} \sim \mathcal{N}(0, 1)$ for neighboring nodes (0 otherwise) and θ_i was set to -1 or $+1$ with equal probability.

We applied PVI to this model, varying the number of particles. To quantify performance, we computed the total variation distance between the inferred marginals and the true marginals. For comparison, we ran a standard mean-field inference algorithm (see Wainwright and Jordan, 2008). Figure 6A shows the results of this analysis, averaged over 100 runs (although the algorithms are in principle deterministic, randomness was introduced by varying the initialization and order of updates). PVI consistently converges to a better approximation than mean-field, even with a single particle. Adding more particles improves the results, but more than 3 particles does not appear to confer any additional improvement.

We can also ask what sort of speed-accuracy tradeoff is achieved by using different numbers of particles: adding more particles increases the time complexity, but also can increase accuracy. To examine this question, we looked at how many iterations of PVI were required before performance exceeded the asymptotic performance of mean-field. We then measured the computation time required to achieve this level of performance with different numbers of particles. As shown in Figure 6B, PVI typically exceeded mean-field after only a few iterations, but this latency was faster for a larger number of particles.

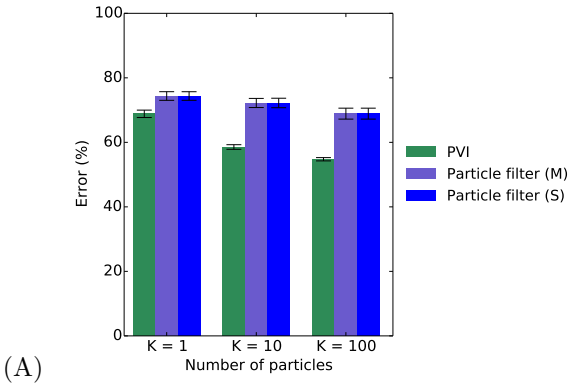

(A)

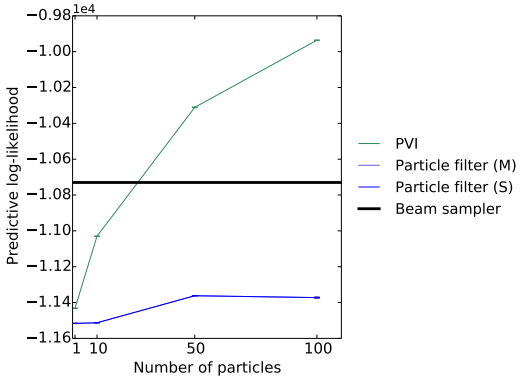

(B)

Figure 5: **Infinite HMM results**. (*A*) Results on 500 synthetic data points generated from an HMM with 10 hidden states. Error is the Hamming distance between the true hidden sequence and the sampled sequence averaged over 50 datasets. M: multinomial resampling; S: stratified resampling. (*B*) Predictive log-likelihood for the "Alice in Wonderland" dataset.

Importantly, using more particles was actually *more* efficient for achieving the same performance level, since fewer iterations were required. This suggests that using a larger number of particles can achieve a favorable speed-accuracy trade-off.

7 Conclusions

This paper introduced a variational framework for particle approximations of discrete probability distributions. We described a practical algorithm for optimizing the approximation, and showed empirically that it can outperform widely-used Monte Carlo algorithms. The key to the success of this approach is an optimal selection of particles: Rather than generating them randomly (as in Monte Carlo algorithms), we deterministically choose a set of unique particles that optimizes the KL divergence between the approximation and the target distribution. Because we are selecting particles optimally, we can achieve good performance with a smaller number of particles compared to Monte Carlo algorithms, thereby improving computational efficiency. Another advantage of PVI is that its deterministic nature eliminates the contribution of Monte Carlo variance to estimation error.

(A)

Number of particles	PVI	SMC
$K = 1$	1.28	1.14s
$K = 10$	3.56s	1.92s
$K = 100$	204.42s	31.99s

(B)

Number of particles	PVI	SMC
$K = 1$	4.73s	1.64s
$K = 10$	41.62s	28.08s
$K = 100$	1685s	211.66s

Table 4: Run time comparison for iHMM. (A) Results using synthetic iHMM dataset from Figure 5A and (B) highlights results obtained by using "Alice in Wonderland" dataset. In case of iHMM, in terms of accuracy, PVI drastically outperforms baseline. When data has long sequences, resampling at every step will produce degeneracy in SMC leading to comparatively lower number of particles with regards to PVI. This leads to PVI having a higher computational time due to increased number of particle. However, the increase in number of particle actually helps PVI achieve accuracy as it better reflects the underlying latent structure of the observations. Theoretically, the only computational overhead for PVI with respect to baseline filtering is due to sorting for picking K best particles.

A consistent problem vexing sequential Monte Carlo methods like particle filtering is the double-edged sword of resampling: this step is necessary to remove conditionally unlikely particles, but the resulting loss of particle diversity can lead to degeneracy. As we showed in our experiments, tuning an ESS threshold for resampling can improve performance, but requires finding a relatively narrow sweet spot for the threshold. PVI achieves comparable performance to the best particle filter by using a deterministic strategy for deleting and replacing particles, avoiding finicky tuning parameters. It is also worth noting two other desirable properties of PVI in this context: (1) the particle set is guaranteed to be diverse because all particles are unique; (2) all the particles have high probability and therefore the propagation of conditionally unlikely particles is

avoided, as happens when particle filtering is run without resampling . We believe that this combination of properties is a key to the superior performance of PVI relative to particle filtering.

Currently, our framework is limited to discrete probability distributions. An important task for future work is to consider how it can be generalized to continuous distributions. One challenge in developing this generalization is that the set of assignments for x_n cannot be enumerated; thus, we require a proposal distribution, raising the question of how to incorporate proposals into the theoretical framework. From a practical perspective, ideas from the Monte Carlo literature may prove useful in identifying good proposal distributions. A related issue is that models with combinatorial latent structure (such as the factorial HMM) may have too many assignments to enumerate completely. It may be desirable to use a proposal distribution in this setting as well.

In summary, PVI harmoniously combines a number of ideas from Monte Carlo and variational methods. The resulting algorithm can achieve performance superior to widely used particle filtering and mean-field methods, though more work is needed to evaluate its performance on a wider range of probabilistic models.

References

Aldous, D. (1985). Exchangeability and related topics. *École d'Été de Probabilités de Saint-Flour XIII1983*, pages 1–198.

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, 98:409.

Antoniak, C. E. (1974). Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *The Annals of Statistics*, pages 1152–1174.

Beal, M. J., Ghahramani, Z., and Rasmussen, C. E. (2002). The infinite hidden Markov model. In *Advances in Neural Information Processing Systems*, pages 577–584.

Besag, J. (1986). On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 259–302.

Daume, H. (2007). Fast search for Dirichlet process mixture models. In *International Conference on Artificial Intelligence and Statistics*, pages 83–90.

Doucet, A., De Freitas, N., Gordon, N., et al. (2001). *Sequential Monte Carlo methods in practice*. Springer New York.

Escobar, M. D. and West, M. (1995). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*, 90:577–588.

Fearnhead, P. (2004). Particle filters for mixture models with an unknown number of components. *Statistics and Computing*, 14(1):11–21.

Flerova, N., Rollon, E., and Dechter, R. (2012). Bucket and mini-bucket schemes for m best solutions over graphical models. In *Graph Structures for Knowledge Representation and Reasoning*, pages 91–118. Springer.

Frank, A., Smyth, P., and Ihler, A. T. (2009). Particle-based variational inference for continuous systems. In *Advances in Neural Information Processing Systems*.

Gershman, S., Hoffman, M., and Blei, D. (2012). Nonparametric variational inference. *Proceedings of the 29th International Conference on Machine Learning*.

Ionides, E. L. (2008). Truncated importance sampling. *Journal of Computational and Graphical Statistics*, 17.

Jaakkola, T. S. and Jordan, M. I. (1998). Improving the mean field approximation via the use of mixture distributions. In Jordan, M., editor, *Learning in Graphical Models*, pages 163–173. Springer.

Lawrence, N. D. (2000). *Variational inference in probabilistic models*. PhD thesis, University of Cambridge.

Quiroga, R. Q., Nadasdy, Z., and Ben-Shaul, Y. (2004). Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural computation*, 16(8):1661–1687.

Robert, C. P. and Casella, G. (2004). *Monte Carlo Statistical Methods*. Springer.

Rosenberg, A. and Hirschberg, J. (2007). V-measure: A conditional entropy-based external cluster evaluation measure. In *EMNLP-CoNLL*, volume 7, pages 410–420.

Teh, Y. W., Jordan, M. I., Beal, M. J., and Blei, D. M. (2006). Hierarchical Dirichlet processes. *Journal of the american statistical association*, 101.

Van Gael, J., Saatci, Y., Teh, Y. W., and Ghahramani, Z. (2008). Beam sampling for the infinite hidden Markov model. In *Proceedings of the 25th international conference on Machine learning*, pages 1088–1095. ACM.

Wainwright, M. J. and Jordan, M. I. (2008). Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1:1–305.

Wang, L. and Dunson, D. B. (2011). Fast Bayesian inference in Dirichlet process mixture models. *Journal of Computational and Graphical Statistics*, 20.

Wood, F. and Black, M. J. (2008). A nonparametric Bayesian alternative to spike sorting. *Journal of Neuroscience Methods*, 173(1):1–12.

Yanover, C. and Weiss, Y. (2003). Finding the M most probable configurations in arbitrary graphical models. In *Advances in Neural Information Processing Systems*, page None.